

database, and document texts. For retrieval from either source, question interpretation maximising the use of general linguistic (and hence world) knowledge, and minimising that of domain-specific knowledge, has both practical and theoretical advantages. The more general a question interpreter is, the more portable it is; and, it is claimed, the better it reproduces human reliance on general knowledge as a support for language use in specific context.

The main challenge for the projects is therefore to provide an adequate general semantics, and to relate general question processing to necessary domain-specific processing effectively. For data questions the problem is translating a meaning representation output by a general analyser into a formal, domain-specific query; for document questions the problem is selecting components of the representation as search terms. A common problem is the treatment of compound nouns.

Both projects make use of an English analyser developed by Boguraev, which uses an ATN parser applying a semantic apparatus derived from Wilks in conjunction with conventional syntax to derive meaning representations in the form of case-labelled dependency trees. Work on the more advanced data question project is currently focussed on the translation component of the question interpreter: this is seen as mapping natural language expressions onto data query language expressions via their respective similar, or related, meaning representation languages. The processing is being done in several stages, applying appropriate sets of production-like rules at each step.

Related research is in progress on the use of a semantic network as an alternative device for linking natural and data languages, and work recently began on further development of case in meaning representations used for translation.

The group's research as a whole seeks to answer the questions:

- a) what are the nature and role of semantic primitives and patterns in text meaning representation
- b) how are primitives and patterns applied in conjunction with syntactic processing in text analysis
- c) how are meaning representations using primitives and patterns related to types of knowledge structure and actual knowledge bases
- d) how are information meaning representations assimilated by associative, inferential, or structure application processes into an existing knowledge store.

#### *References*

- B.K. Boguraev, Automatic resolution of linguistic ambiguities. Ph.D. Thesis, University of Cambridge, 1979; Technical Report 11, Computer Laboratory, University of Cambridge, 1979.
- B.K. Boguraev and K. Sparck Jones, "A general semantic analyser for database access", Proceedings of the Seventh International Joint Conference on Artificial Intelligence, 1981, 443-445.

B.K. Boguraev and K. Sparck Jones, "A natural language analyser for database access", Information Technology: Research and Development, 1, 1982 (in press).

A.W.S. Cater, "Analysing English text: a nondeterministic approach with limited memory", Proceedings of the AISB-80 Conference on Artificial Intelligence, 1980.

---

### **Robust Man-Machine Interfaces and Dialog Modelling Carnegie-Mellon University**

Mark Boggs, Jaime Carbonell, Jr., Robert Frederking,  
Phil Hayes, George Mouradian, Don Kosy,  
Michael Mauldin, Hiromichi Fujisawa

Department of Computer Science  
Carnegie-Mellon University  
Schenley Park  
Pittsburgh, PA 15217

A number of projects at the Carnegie-Mellon University Computer Science Department address issues in Natural Language Processing. Since several of these projects share researchers and a similar view of the world, we have listed them in a single summary. Our general interests are in robust man-machine interfaces and the modelling of human dialogs.

#### *Parsing for Limited-Domain Natural Language Interfaces* (Boggs, Carbonell, Frederking, Hayes, Mouradian)

We are interested in providing natural language access to interactive computer systems of various kinds, including data base query and update, monitor-level command interpreters, statistical data analysis packages, or any other domain in which actions, relations, and objects can be subject to unambiguous semantic categorization. These restrictions allow us to build robust interfaces (by exploiting the underlying domain semantics) that can efficiently parse well-formed utterances, and can cope with a substantial degree of grammatical deviation.

More specifically, we are focusing our investigations on the following:

- Extragrammatical input: Whenever people use language spontaneously, they tend to deviate from standard grammar either through mistakes, or by economy of expression manifested as ellipsis, abbreviations, omitted articles or prepositions, and other such devices. It is undesirable for an interactive interface to reject such extragrammatical input, so we are developing techniques to parse

ellipsed and ungrammatical input flexibly, i.e. to correct errors uniquely or to within a small set of possible interpretations. This process requires us to interpret as much of the input as possible even when undecipherable segments must be skipped.

- Application of Domain Semantics: Because of their considerable power of selectivity, the strong semantic typing assumptions mentioned above make natural language processing in restricted domains quite tractable, and essentially all our work in applied natural language processing has relied on this fact. We are investigating techniques for applying these selectional restrictions to make the parsing of both grammatical and ungrammatical input highly efficient.

Recently, we have been investigating an approach to parsing based on the use of several distinct parsing strategies, selected among on a dynamic basis, with each strategy tailored to a specific type of language construction. This multi-strategy, construction-specific approach to parsing domain-specific natural language appears to offer considerable advantages over our earlier more uniform approach [15] for the two goals itemized above (see [5, 11, 12, 7]), as well as for the recognition and localization of ambiguity and the use of domain-oriented language definitions [13]. The approach has also proved useful in the flexible parsing [1] of artificial languages of the type commonly found in interactive command interfaces.

Thus far, we have built four pilot parsers and are experimenting with a more comprehensive multi-strategy language analyzer.

- FLEXP - flexible, bottom-up, single strategy, pattern-matching parser (applied to a message system domain).

- CASPAR - two-strategy experimental parser (case-frame instantiation and directed pattern matching).

- DYPAR - three-strategy experimental parser (semantic grammars, partial pattern matching, and linguistic transformations) being adapted as a "stopgap" front end to several systems including the R1/XSEL expert system and a factory scheduling system.

- COUSIN - a friendly interactive operating system command interface; the parser for this interface parses an artificial command language, using the same flexible techniques as we have developed for restricted natural language interfaces.

- MULTIPAR - Everything we always wanted in a robust, transportable, task-oriented, multi-strategy natural language interface but were afraid to implement. (Well, not quite -- we are working on it.) MULTIPAR should replace DYPAR as a more effective "real-world applications" parser.

#### *Natural Language Functionality without Cognitive Simulation*

(Hayes)

Much of the power and convenience of natural language comes from the use of such devices as anaphora, ellipsis, and implicit context

switching. These natural communication mechanisms allow much redundant information to be omitted without impairing understanding. However, as research into these phenomena has shown, human performance using these mechanisms can require considerable amounts of deep inferencing and cognitive modelling, much more, in particular, than can currently be performed by an interactive interface requiring real-time response. The utility of these mechanisms nevertheless makes it desirable to provide them in practical natural language interfaces, and we are, therefore, interested in the development of substitute mechanisms that provide a level of functionality equivalent or similar to human performance but without the computationally expensive inferencing and cognitive modelling.

We have already produced such a mechanism for simple and set selection pronoun anaphora [14]. The mechanism approximates human performance wherever this can be done without cognitive simulation. When this is impossible, the mechanism behaves radically differently from human performance, but according to simple to understand rules. In this way, an interface user can either rely on his experience of human performance or his knowledge of the artificial, but simple, substitute rules to predict the behaviour of the system to his inputs, and thus construct his inputs to use the mechanism to best advantage. We are planning to use a similar approach to the provision of other natural communication mechanisms in practical natural language interfaces.

#### *Metaphor Interpretation*

(Carbonell)

Metaphors abound in natural language text, from narratives to newspaper articles and textbooks. The present investigation focuses on interpreting creative (i.e., non-frozen) metaphors in context, rather than attempting to circumvent them or treat them as a deviant use of language. Two paths have been followed:

- We established a small set of very general metaphor mappings ubiquitous in everyday language. When metaphors are encountered, we attempt to select the appropriate general mapping, and interpret the surrounding text as an instantiation of that mapping. In essence, this provides a top-down component constraining inferences in metaphor comprehension. We find that well over half of the metaphors encountered in common language can be classified as instances of some 40 to 50 generalized mappings [6, 4].

- We analyze the reason why a given creative metaphor is used (vs. a literal rendition or a different metaphor) by extracting the knowledge implicitly conveyed in the underlying analogical mapping from the source domain to the less-understood target domain. Metaphors are generated with a purpose -- often to enrich the knowledge available to the understander. Present investigations indicate that certain categories of knowledge (such as goals and planning strategies) present in the source domain are typically mapped invariant into the target domain, whereas other categories (such as physical descriptors of objects) are lost in the mapping process [8].

*Dialog Modelling*

(Frederking, Carbonell)

Our dialog modelling project addresses issues of subjective comprehension in conversations, topic transitions, focus of attention, and extracting rudimentary models of the social interaction underlying discourse. As such, our objectives differ somewhat from the more formal, but less cognitive approach of speech-act categorization. We have implemented a pilot dialog model called MICS (Mixed-Initiative Conversation System) that operates by focusing attention on the inferred conversational goals of the participants. MICS embodies a set of conversational principles (similar to, but more detailed than Grice's Maxims) that in conjunction with the goal-focusing mechanism enable it to maintain conversational coherence when changing, elaborating, terminating or digressing from the dialog topic [2, 9].

More recently, we have been developing a rule-based system which attempts to capture the notion of "topics" using dynamically generated and modified data structures. This system simulates a conversation participant in a non-task oriented conversation, and attempts to use general principles of conversational interest based on the goals of the speaker, listener, and person being discussed to determine what to say next. It brings scripts, dynamically changing goal trees, and a semantic network to bear on this problem [10].

*Learning from Reading: Acquiring Knowledge from Natural Language Textbook Input*

(Kosy)

This study concerns a variant of learning by being told which might be called learning by being taught. To teach a person a new concept, it is insufficient to simply pair the name of the concept with its definition if one wishes the person to remember it. Instead, one must provide motivation for the concept, explanations of its origin, its relation to other concepts, examples of its use, and exercises by which the person can practice identifying instances and drawing inferences based on the concepts identified.

We have been investigating the kinds of expository prose that textbook authors use to provide the information a student needs. Research has centered on a particular case involving the basic concepts and vocabulary of accounting as they are described in the first chapter of a freshman accounting text. Since this is "real text", the parsing, interpretation, and assimilation of its content has posed a real challenge to existing NL techniques. For example, the term assets is introduced by the sentence "The term 'assets' takes its place." To determine what this term might refer to, a reader must

- determine the referent of "it",
- determine the "place" of it in this context,
- determine what it means for a new term to take that place,
- and make the necessary additions to his knowledge to incorporate the new term and its meaning.

To meet the challenge of this and other sentences that are similarly underspecified out of context, we have hypothesized a system that employs Sidner's methods of anaphora resolution coupled to a base of lexical knowledge and knowledge of rhetorical patterns. Together, these are used to assemble new assertions to go into a discourse database. Of central concern has been how to specify and represent the knowledge in a way that can serve as both a source for solving comprehension problems and as a sink for new information derived from what the input text asserts. Work is proceeding on development of a suitable knowledge base for the several domains the accounting text makes use of and on a procedural specification of how an interpreter of segments produced by a wait-and-see style parser should use it.

*Other Projects*

In addition to the projects discussed above, other investigations are in progress. Only cursory mention is made of these investigations either because they are only in their nascent stages or because they are presently on the proverbial back burner.

- Inferring the definitions of unknown words from context. This project was undertaken in the context of the POLITICS system, but is not presently being pursued. [3] Carbonell

- Generating explanations of expert-system behavior by analyzing dynamic traces as well as static rules and data. This is a new project, where much of the emphasis is directed at determining how a model of the user affects what a system should say, and how its explanations should be phrased. Mauldin, Carbonell

- Conversing with the users of the previously mentioned factory scheduling system. This is also a new project, with its main goals being the implementation of a case-frame parser in a true production system and the integration of case-frame parsing, language generation, and dialog rules within a single production system. Frederking

- A re-implementation of Hendrix's LIFER parser, including a user-extensible grammar, as a component of a speech-understanding image processing system. This component of the system is finished, and its author has returned to Japan. Fujisawa

*References*

- [1] Ball, J. E., Hayes, P. J., and Szekely, P. A. Design for COUSIN: A Cooperative User Interface. Technical Report, Carnegie-Mellon University, Computer Science Department, 1981.
- [2] Carbonell, J. G. Intentionality and Human Conversation. In Proceedings of the Second Conference on Theoretical Issues in Natural Lang. Proc., pages 141-148. TINLAP-II, 1978.
- [3] Carbonell, J. G. Towards a Self-Extending Parser. In Proceedings of the 17th Meeting of the Association for Computational Linguistics, pages 3-7. ACL-79, 1979.

- [4] Carbonell, J. G. Metaphor - A Key to Extensible Semantic Analysis. In Proceedings of the 18th Meeting of the Association for Computational Linguistics. ACL-80, 1980.
  - [5] Carbonell, J. G. and Hayes, P. J. Dynamic Strategy Selection in Flexible Parsing. In Proc. of 19th Annual Meeting of the Assoc. for Comput. Ling.. Stanford University, June, 1981.
  - [6] Carbonell, J. G. Metaphor Comprehension. In W. Lehnert and M. Ringle (editors), Knowledge Representation for Language Processing Systems, . New Jersey: Erlbaum, 1981.
  - [7] Carbonell, J. G. Towards a Robust, Task-Oriented Natural Language Interface. In A. N. Badler (editor), Workshop/Symposium on Human Computer Interaction. Georgia Tech Information Sciences, March, 1981.
  - [8] Carbonell, J. G. Invariance Hierarchies in Metaphor Interpretation. In Proceedings of the Third Meeting of the Cognitive Science Society, pages 292-295. CognitiveScienceSoc.", August, 1981.
  - [9] Carbonell, J. G. Subjective Understanding: Computer Models of Belief Systems. Ann Arbor, MI: UMI research press, 1981.
  - [10] Frederking, R. E. A Rule-based Conversation Participant. In Proceedings of the 19th Meeting of the Association for Computational Linguistics. ACL-81, 1981.
  - [11] Hayes, P. J. and Carbonell, J. G. Multi-Strategy Parsing and its Role in Robust Man-Machine Communication. Technical Report, Carnegie-Mellon University, Computer Science Department, May, 1981.
  - [12] Hayes, P. J. and Carbonell, J. G. Multi-Strategy Construction-Specific Parsing for Flexible Data Base Query and Update. In Proc. Seventh Int. Jt. Conf. on Artificial Intelligence, pages 432-439. Univ. of British Columbia, Vancouver, August 1981.
  - [13] Hayes P. J. Focused Interaction in Flexible Parsing. In Proc. of 19th Annual Meeting of the Assoc. for Comput. Ling.. Stanford University, June, 1981.
  - [14] Hayes, P. J. Anaphora in Limited Domain Systems. In Proc. Seventh Int. Jt. Conf. on Artificial Intelligence, pages 416- 422. Vancouver, 1981.
  - [15] Hayes, P. J. and Mouradian, G. V. Flexible Parsing. In Proc. of 18th Annual Meeting of the Assoc. for Comput. Ling., pages 97-103. Philadelphia, June, 1980.
- 

## Center for the Study of Reading (BBN Branch)

Bertram Bruce  
Center for the Study of Reading  
Bolt, Beranek and Newman branch

The Center for the Study of Reading (CSR) is a joint project of the University of Illinois at Urbana-Champaign and Bolt Beranek and Newman Inc. It was established in 1976 to study the problems of learning to read, with a special emphasis on the difficulties many children have with reading comprehension. The CSR draws together researchers from a variety of disciplines including artificial intelligence, cognitive psychology, linguistics, anthropology, and education.

By far the largest part of the CSR is at the University of Illinois, where Richard C. Anderson and Tom Anderson are Co-Directors. Technical reports of the CSR are published there also; a list may be obtained by writing to Center for the Study of Reading, 51 Gerty Drive, University of Illinois at Urbana-Champaign, Champaign, Illinois 61820, USA. The summary given below deals only with the part of the CSR located at Bolt Beranek and Newman Inc.

Researchers at BBN currently working on the project include Marilyn Adams, Bertram Bruce (Principal Investigator), Allan Collins, John Frederiksen, Dedre Gentner, Andee Rubin, Ed Smith, Kathleen Starr and Cindy Steinberg. Phil Cohen at Oregon State University, Denis Newman at U.C.S.D., Martin Ringle at Vassar, Mark Seidenberg at McGill, and Robert Tierney at Harvard are also closely affiliated with the BBN group. In the past, Bonnie Webber, Bill Woods, Kathy Larkin, Bill Huggins, Rachel Reichman, Marilyn Shatz, John Brown, Aleida Inglis, Scott Fertig and others have had significant involvements with the research.

Research of the BBN group has covered a range of topics too diverse to describe in any detail here. Our progress report listed below is available for those who would like more information. Among the most active areas of current research of the BBN group are the following:

1. Computer activities and tools for teaching reading and writing.
2. Models of interacting plans for use in analyzing stories and the comprehension of stories.
3. Communication in different modalities, e.g., lecture vs. conversation, teletype vs. telephone.
4. Plausible reasoning, e.g., if A were true I would have known it, I don't know it, therefore it's not true.
5. Conceptual parsing--how we divide up the world into objects and relations.
6. Scientific vs. literary metaphors.
7. The author-reader relationship, including the explicit and implicit embedding of stories within stories.